

Two-Dimensional Pattern-Coupled Sparse Bayesian Learning via Generalized Approximate Message Passing

Jun Fang, Lizao Zhang, and Hongbin Li, *Senior Member, IEEE*

Abstract—We consider the problem of recovering two-dimensional (2-D) block-sparse signals with *unknown* cluster patterns. Two-dimensional block-sparse patterns arise naturally in many practical applications such as foreground detection and inverse synthetic aperture radar imaging. To exploit the block-sparse structure, we introduce a 2-D pattern-coupled hierarchical Gaussian prior model to characterize the statistical pattern dependencies among neighboring coefficients. Unlike the conventional hierarchical Gaussian prior model where each coefficient is associated independently with a unique hyperparameter, the pattern-coupled prior for each coefficient not only involves its own hyperparameter, but also its immediate neighboring hyperparameters. Thus the sparsity patterns of neighboring coefficients are related to each other and the hierarchical model has the potential to encourage 2-D structured-sparse solutions. An expectation-maximization (EM) strategy is employed to obtain the maximum a posteriori (MAP) estimate of the hyperparameters, along with the posterior distribution of the sparse signal. In addition, the generalized approximate message passing (GAMP) algorithm is embedded into the EM framework to efficiently compute an approximation of the posterior distribution of hidden variables, which results in a significant reduction in computational complexity. Numerical results are provided to illustrate the effectiveness of the proposed algorithm.

Index Terms—Pattern-coupled sparse Bayesian learning, block-sparse structure, expectation-maximization (EM), generalized approximate message passing (GAMP).

I. INTRODUCTION

Compressed sensing is a recently emerged technique for signal sampling and data acquisition which enables to recover sparse signals from undersampled linear measurements

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{M \times N}$ is a sampling matrix with $M \ll N$, \mathbf{x} denotes an N -dimensional sparse signal, and \mathbf{w} denotes the additive noise. The problem has been extensively studied and a variety of algorithms, e.g. the orthogonal matching pursuit (OMP) algorithm [1], the basis pursuit (BP) method [2], the iterative reweighted ℓ_1 and ℓ_2 algorithms [3], and the sparse Bayesian learning method [4]–[6] were proposed. In

many practical applications, in addition to the sparse structure, sparse signals may exhibit two-dimensional cluster patterns that can be utilized to enhance the recovery performance. For example, the target of interest in the synthetic aperture radar/inverse synthetic aperture radar (SAR/ISAR) images often demonstrates continuity in both the range and cross-range domains [7]. In video surveillance, the foreground image exhibits a cluster pattern since the foreground objects (humans, cars, text etc.) generally occupy a small continuous region of the scene [8]. Besides these, block-sparsity is also present in temporal observations of a time-varying block-sparse signal whose support varies slowly over time [9].

Analyses [10]–[12] show that exploiting the inherent block-sparse structure not only leads to relaxed conditions for exact reconstruction, but also helps improve the recovery performance considerably. A number of algorithms have been proposed for recovering block-sparse signals over the past few years, e.g., block-OMP [12], mixed ℓ_2/ℓ_1 norm-minimization [10], group LASSO [13], model-based CoSaMP [11], and block-sparse Bayesian learning [14], [15]. These algorithms, however, require *a priori* knowledge of the block partition (e.g. the number of blocks and location of each block) such that the coefficients in each block are grouped together and enforced to share a common sparsity pattern. In practice, the prior information about the block partition of sparse signals is often unavailable, especially for two-dimensional signals since the block partition of a two-dimensional signal involves not only the location but also the shape of each block. For example, foreground images have irregular and unpredictable cluster patterns which are very difficult to be estimated *a priori*. To address this difficulty, a few sophisticated Bayesian methods which do not need the knowledge of the block partition were developed. In [16], a “spike-and-slab” prior model was proposed, where by introducing dependencies among mixing weights, the prior model has the potential to encourage sparsity and promote a tree structure simultaneously. This “spike-and-slab” prior model was later extended to accommodate block-sparse signals [7], [17]. Nevertheless, for the “spike-and-slab” prior introduced in [16], [17], the posterior distribution cannot be derived analytically, and a Markov chain Monte Carlo (MCMC) sampling method has to be employed for Bayesian inference. In [18], [19], a graphical prior, also referred to as the “Boltzmann machine”, is employed as a prior on the sparsity support in order to induce statistical dependencies between atoms. With such a prior, the maximum a posteriori (MAP) estimator requires an exhaustive search over all possible sparsity

Jun Fang, and Lizao Zhang are with the National Key Laboratory of Science and Technology on Communications, University of Electronic Science and Technology of China, Chengdu 611731, China, Email: JunFang@uestc.edu.cn

Hongbin Li is with the Department of Electrical and Computer Engineering, Stevens Institute of Technology, Hoboken, NJ 07030, USA, E-mail: Hongbin.Li@stevens.edu

This work was supported in part by the National Science Foundation of China under Grant 61172114, and the National Science Foundation under Grant ECCS-1408182.

patterns. To overcome the intractability of the combinatorial search, a greedy method [18] and a variational mean-field approximation method [19] were developed to approximate the MAP. In [20], to cope with the unknown cluster pattern, an expanded model is employed by assuming that the original sparse signal is a superposition of a number of overlapping blocks, and the coefficients in each block share the same sparsity pattern. Conventional block sparse Bayesian learning algorithms such as those in [15] can then be applied to the expanded model.

Recently in [21], we proposed a pattern-coupled hierarchical Gaussian prior model to exploit the unknown block-sparse structure. Unlike the conventional hierarchical Gaussian prior model [4]–[6] where each coefficient is associated independently with a unique hyperparameter, the pattern-coupled prior for each coefficient not only involves its own hyperparameter, but also its immediate neighboring hyperparameters. This pattern-coupled hierarchical model is effective and flexible to capture any underlying block-sparse structures, without requiring the prior knowledge of the block partition. Numerical results show that the pattern-coupled sparse Bayesian learning (PC-SBL) method renders competitive performance for block-sparse signal recovery. Nevertheless, a major drawback of the method is that it requires computing an $N \times N$ matrix inverse at each iteration, and thus has a cubic complexity in terms of the signal dimension. This high computational cost prohibits its application to problems with even moderate dimensions. Also, [21] only considers recovery of one-dimensional block-sparse signals. In this paper, we generalize the pattern-coupled hierarchical model to the two-dimensional (2-D) scenario in order to leverage block-sparse patterns arising from 2-D sparse signals. To address the computational issue, we resort to the generalized approximate message passing (GAMP) technique [22] and develop a computationally efficient method. Specifically, the algorithm is developed within an expectation-maximization (EM) framework, using the GAMP to efficiently compute an approximation of the posterior distribution of hidden variables. The hyperparameters associated with the hierarchical Gaussian prior are learned by iteratively maximizing the Q-function which is calculated based on the posterior approximation obtained from the GAMP. Simulation results show that the proposed method presents superior recovery performance for block-sparse signals, meanwhile achieving a significant reduction in computational complexity.

The rest of the paper is organized as follows. In Section II, we introduce a 2-D pattern coupled hierarchical Gaussian framework to model the sparse prior and the pattern dependencies among the neighboring coefficients. In Section III, a GAMP-based EM algorithm is developed to obtain the maximum a posteriori (MAP) estimate of the hyperparameters, along with the posterior distribution of the sparse signal. Simulation results are provided in Section IV, followed by concluding remarks in Section V.

II. BAYESIAN MODEL

We consider the problem of recovering a two-dimensional block-sparse signal $\mathbf{X} \in \mathbb{R}^{Q \times L}$ from compressed noisy

measurements

$$\mathbf{y} = \mathbf{f}(\mathbf{X}) + \mathbf{w} \quad (2)$$

where $\mathbf{y} \in \mathbb{R}^M$ denotes the compressed measurement vector, $\mathbf{f}(\cdot)$ is a linear map: $\mathbb{R}^{Q \times L} \rightarrow \mathbb{R}^M$, with $M \ll N \triangleq QL$, and $\mathbf{w} \in \mathbb{R}^M$ is an additive multivariate Gaussian noise with zero mean and covariance matrix $\sigma^2 \mathbf{I}$. Let $\mathbf{x} \triangleq \text{vec}(\mathbf{X})$, the linear map $\mathbf{f}(\mathbf{X})$ can generally be expressed as

$$\mathbf{f}(\mathbf{X}) = \mathbf{A}\mathbf{x} \quad (3)$$

where $\mathbf{A} \in \mathbb{R}^{M \times N}$ denotes the measurement matrix. In the special case where $\mathbf{f}(\mathbf{X}) = \text{vec}(\mathbf{B}\mathbf{X})$, then we have $\mathbf{A} \triangleq \mathbf{I} \otimes \mathbf{B}$, in which \otimes stands for the Kronecker product. The above model (2) arises in image applications where signals are multi-dimensional in nature, or in the scenario where multiple snapshots of a time-varying sparse signal are available. In these applications, signals usually exhibit two-dimensional cluster patterns that can be utilized to improve the recovery accuracy. To leverage the underlying block-sparse structures, we introduce a 2-D pattern-coupled Gaussian prior model which is a generalization of our previous work [21]. Before proceeding, we provide a brief review of the conventional hierarchical Gaussian prior model [4], and some of its extensions.

A. Review of Conventional Gaussian Prior Model

For ease of exposition, we consider the prior model for the two-dimensional signal \mathbf{X} instead of its one-dimensional form \mathbf{x} . Let $x_{q,l}$ denote the (q,l) th entry of \mathbf{X} . In the conventional sparse Bayesian learning framework [4], a two-layer hierarchical Gaussian prior was employed to promote the sparsity of the solution. In the first layer, coefficients $\{x_{q,l}\}$ of \mathbf{X} are assigned a Gaussian prior distribution

$$p(\mathbf{X}|\boldsymbol{\alpha}) = \prod_{q=1}^Q \prod_{l=1}^L \mathcal{N}(x_{q,l}|0, \alpha_{q,l}^{-1}) \quad (4)$$

where $\alpha_{q,l}$ is a non-negative hyperparameter controlling the sparsity of the coefficient $x_{q,l}$. The second layer specifies Gamma distributions as hyperpriors over the hyperparameters $\boldsymbol{\alpha} \triangleq \{\alpha_{q,l}\}$, i.e.

$$p(\boldsymbol{\alpha}) = \prod_{q=1}^Q \prod_{l=1}^L \text{Gamma}(\alpha_{q,l}|a, b) \quad (5)$$

As discussed in [4], for properly chosen a and b , this hyperprior allows the posterior mean of $\alpha_{q,l}$ to become arbitrarily large. As a consequence, the associated coefficient $x_{q,l}$ will be driven to zero, thus yielding a sparse solution. This conventional hierarchical model, however, does not encourage structured-sparse solutions since the sparsity of each coefficient is determined by its own hyperparameter and the hyperparameters are independent of each other. In [14], [15], the above hierarchical model was generalized to deal with block-sparse signals, in which a group of coefficients sharing the same sparsity pattern are assigned a multivariate Gaussian prior parameterized by a common hyperparameter. Nevertheless, this model requires the knowledge of the block partition to determine which coefficients should be grouped and assigned a common hyperparameter.

B. Proposed 2-D Pattern-Coupled Hierarchical Model

To exploit the 2-D block-sparse structure, we utilize the fact that the sparsity patterns of neighboring coefficients are statistically dependent. To capture the pattern dependencies among neighboring coefficients, the Gaussian prior for each coefficient $x_{q,l}$ not only involves its own hyperparameter $\alpha_{q,l}$, but also its immediate neighbor hyperparameters. Specifically, a prior over \mathbf{X} is given by

$$p(\mathbf{X}|\boldsymbol{\alpha}) = \prod_{q=1}^Q \prod_{l=1}^L \mathcal{N}(x_{q,l}|0, \delta_{q,l}^{-1}) \quad (6)$$

where

$$\delta_{q,l} \triangleq \alpha_{q,l} + \beta \sum_{(i,j) \in N_{(q,l)}} \alpha_{i,j} \quad (7)$$

in which $N_{(q,l)}$ denotes the neighborhood of the grid point (q, l) , i.e. $N_{(q,l)} \triangleq \{(q, l-1), (q, l+1), (q-1, l), (q+1, l)\}^1$, and $\beta \in [0, 1]$ is a parameter indicating the pattern relevance between the coefficient $x_{q,l}$ and its neighboring coefficients. Clearly, this model is an extension of our previous prior model [21] to the two-dimensional case. When $\beta = 0$, the prior model (6) reduces to the conventional sparse Bayesian learning model. When $\beta > 0$, we see that the sparsity of each coefficient $x_{q,l}$ is not only controlled by the hyperparameter $\alpha_{q,l}$, but also by the neighboring hyperparameters $S_{\alpha_{q,l}} \triangleq \{\alpha_{i,j} | (i,j) \in N_{(q,l)}\}$. The coefficient $x_{q,l}$ will be driven to zero if $\alpha_{q,l}$ or any of its neighboring hyperparameters goes to infinity. In other words, suppose $\alpha_{q,l}$ approaches infinity, then not only its corresponding coefficient $x_{q,l}$ will be driven to zero, the neighboring coefficients $S_{x_{q,l}} \triangleq \{x_{i,j} | (i,j) \in N_{(q,l)}\}$ will decrease to zero as well. We see that the sparsity patterns of neighboring coefficients are related to each other through their shared hyperparameters. On the other hand, for any pair of neighboring coefficients, each of them has its own hyperparameters that are not shared by the other coefficient. Hence, no coefficients are pre-specified to share a common sparsity pattern, which enables the prior to provide flexibility to model any block-sparse structures.

Following [4], we use Gamma distributions as hyperpriors over the hyperparameters $\{\alpha_{q,l}\}$, i.e.

$$p(\boldsymbol{\alpha}) = \prod_{q=1}^Q \prod_{l=1}^L \text{Gamma}(\alpha_{q,l} | a, b) \quad (8)$$

where we set $a > 1$, and $b = 10^{-6}$. The choice of a will be elaborated later in our paper. Also, the noise variance $\sigma^2 \triangleq 1/\gamma$ is assumed unknown, and to estimate this parameter, we place a Gamma hyperprior over γ , i.e.

$$p(\gamma) = \text{Gamma}(\gamma | c, d) \quad (9)$$

where we set $c = 1$ and $d = 10^{-6}$.

¹Note that for edge grid points, they only have two or three immediate neighboring points, in which case the definition of $N_{(q,l)}$ changes accordingly.

III. PROPOSED ALGORITHM

We now proceed to perform Bayesian inference for the proposed pattern-coupled hierarchical model. The following model is considered since the linear map $\mathbf{f}(\mathbf{X})$ can be expressed as $\mathbf{f}(\mathbf{X}) = \mathbf{A}\mathbf{x}$

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w} \quad (10)$$

We first translate the prior for the two-dimensional signal \mathbf{X} to a prior for its one-dimensional form \mathbf{x} . From (6), the prior over \mathbf{x} can be expressed as

$$p(\mathbf{x}|\boldsymbol{\alpha}) = \prod_{n=1}^N \mathcal{N}(x_n | 0, \eta_n^{-1}) \quad (11)$$

where

$$\eta_n \triangleq \alpha_n + \beta \sum_{i \in N_{(n)}} \alpha_i \quad (12)$$

in which $N_{(n)}$ denotes the neighbors of the point (q, l) on the two-dimensional grid, i.e. $N_{(n)} \triangleq \{(l-2)Q + q, lQ + q, (l-1)Q + q - 1, (l-1)Q + q + 1\}^2$. The relation between n and (q, l) is given by $n = (l-1)Q + q$, that is, $l = \lceil n/Q \rceil$, and $q = n \bmod Q$, in which $\lceil \cdot \rceil$ denotes the ceiling operator. Note that for notational convenience, we, with a slight abuse of notation, use x_n to denote the n th entry of \mathbf{x} and α_n to denote the hyperparameter associated with the coefficient x_n . Also, let $\boldsymbol{\alpha} \triangleq \{\alpha_n\}$ since its exact meaning remains unaltered. From (8), we have

$$p(\boldsymbol{\alpha}) = \prod_{n=1}^N \text{Gamma}(\alpha_n | a, b) \quad (13)$$

An expectation-maximization (EM) algorithm can be developed for learning the sparse signal \mathbf{x} as well as the hyperparameters $\{\boldsymbol{\alpha}, \gamma\}$. In the EM formulation, the signal \mathbf{x} is treated as hidden variables, and we iteratively maximize a lower bound on the posterior probability $p(\boldsymbol{\alpha}, \gamma | \mathbf{y})$ (this lower bound is also referred to as the Q-function). Briefly speaking, the algorithm alternates between an E-step and a M-step. In the E-step, we need to compute the posterior distribution of \mathbf{x} conditioned on the observed data and the hyperparameters estimated from the t th iteration, i.e.

$$p(\mathbf{x} | \mathbf{y}, \boldsymbol{\alpha}^{(t)}, \gamma^{(t)}) \propto p(\mathbf{x} | \boldsymbol{\alpha}^{(t)}) p(\mathbf{y} | \mathbf{x}, \gamma^{(t)}) \quad (14)$$

It can be readily verified that the posterior $p(\mathbf{x} | \mathbf{y}, \boldsymbol{\alpha}^{(t)}, \gamma^{(t)})$ follows a Gaussian distribution with its mean and covariance matrix given respectively by

$$\boldsymbol{\mu} = \gamma^{(t)} \boldsymbol{\Phi} \mathbf{A}^T \mathbf{y} \quad (15)$$

$$\boldsymbol{\Phi} = (\gamma^{(t)} \mathbf{A}^T \mathbf{A} + \mathbf{D})^{-1} \quad (16)$$

where $\mathbf{D} \triangleq \text{diag}(\eta_1^{(t)}, \dots, \eta_N^{(t)})$. The Q-function, i.e. $E_{\mathbf{x} | \mathbf{y}, \boldsymbol{\alpha}^{(t)}, \gamma^{(t)}}[\log p(\boldsymbol{\alpha}, \gamma | \mathbf{y})]$, can then be computed, where the operator $E_{\mathbf{x} | \mathbf{y}, \boldsymbol{\alpha}^{(t)}, \gamma^{(t)}}[\cdot]$ denotes the expectation with respect to the posterior distribution $p(\mathbf{x} | \mathbf{y}, \boldsymbol{\alpha}^{(t)}, \gamma^{(t)})$. In the M-step,

²For edge grid points which have only two or three neighboring points, the definition of $N_{(n)}$ changes accordingly.

we maximize the Q-function with respect to the hyperparameters $\{\alpha, \gamma\}$.

It can be seen that the EM algorithm, at each iteration, requires to update the posterior distribution $p(\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)})$, which involves computing an $N \times N$ matrix inverse. Thus the EM-based algorithm has a computational complexity of $\mathcal{O}(N^3)$ flops, and therefore is not suitable for many real-world applications involving large dimensions. In the following, we will develop a computationally efficient algorithm by resorting to the generalized approximate message passing (GAMP) technique [22]. GAMP is a very-low-complexity Bayesian iterative technique recently developed [22] for obtaining an approximation of the posterior distribution $p(\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)})$. It therefore can naturally be embedded within the EM framework to replace the computation of the true posterior distribution. From GAMP's point of view, the hyperparameters $\{\alpha, \gamma\}$ are considered as known. The hyperparameters can be updated in the M-step based on the approximate posterior distribution of \mathbf{x} . We now proceed to derive the GAMP algorithm for the pattern-coupled Gaussian hierarchical prior model.

A. Pattern-Coupled Hierarchical Gaussian GAMP

GAMP was developed in a message passing-based framework. By using central-limit-theorem approximations, message passing between variable nodes and factor nodes can be greatly simplified, and the loopy belief propagation on the underlying factor graph can be efficiently performed. As noted in [22], [23], the central-limit-theorem approximations become exact in the large-system limit under i.i.d. zero-mean sub-Gaussian \mathbf{A} .

For notational convenience, let $\boldsymbol{\theta} \triangleq \{\alpha, \gamma\}$ denote the hyperparameters. Firstly, GAMP assumes posterior independence among hidden variables $\{x_n\}$ and approximates the true posterior distribution $p(x_n|\mathbf{y}, \boldsymbol{\theta})$ by

$$\hat{p}(x_n|\mathbf{y}, \hat{r}_n, \tau_n^r, \boldsymbol{\theta}) = \frac{p(x_n|\boldsymbol{\theta})\mathcal{N}(x_n|\hat{r}_n, \tau_n^r)}{\int_x p(x_n|\boldsymbol{\theta})\mathcal{N}(x_n|\hat{r}_n, \tau_n^r)} \quad (17)$$

where \hat{r}_n and τ_n^r are quantities iteratively updated during the iterative process of the GAMP algorithm. Here, we have dropped their explicit dependence on the iteration number k for simplicity. Substituting (11) into (17), it can be easily verified that the approximate posterior $\hat{p}(x_n|\mathbf{y}, \hat{r}_n, \tau_n^r, \boldsymbol{\theta})$ follows a Gaussian distribution with its mean and variance given respectively as

$$\mu_n^x \triangleq \frac{\hat{r}_n}{1 + \eta_n \tau_n^r} \quad (18)$$

$$\phi_n^x \triangleq \frac{\tau_n^r}{1 + \eta_n \tau_n^r} \quad (19)$$

Another approximation is made to the noiseless output $z_m \triangleq \mathbf{a}_m^T \mathbf{x}$, where \mathbf{a}_m^T denotes the m th row of \mathbf{A} . GAMP approximates the true marginal posterior $p(z_m|\mathbf{y}, \boldsymbol{\theta})$ by

$$\hat{p}(z_m|\mathbf{y}, \hat{p}_m, \tau_m^p, \boldsymbol{\theta}) = \frac{p(y_m|z_m, \boldsymbol{\theta})\mathcal{N}(z_m|\hat{p}_m, \tau_m^p)}{\int_z p(y_m|z_m, \boldsymbol{\theta})\mathcal{N}(z_m|\hat{p}_m, \tau_m^p)} \quad (20)$$

where \hat{p}_m and τ_m^p are quantities iteratively updated during the iterative process of the GAMP algorithm. Again, here we

dropped their explicit dependence on the iteration number k . Under the additive white Gaussian noise assumption, we have $p(y_m|z_m, \boldsymbol{\theta}) = \mathcal{N}(y_m|z_m, 1/\gamma)$. Thus $\hat{p}(z_m|\mathbf{y}, \hat{p}_m, \tau_m^p, \boldsymbol{\theta})$ also follows a Gaussian distribution with its mean and variance given by

$$\mu_m^z \triangleq \frac{\tau_m^p \gamma y_m + \hat{p}_m}{1 + \gamma \tau_m^p} \quad (21)$$

$$\phi_m^z \triangleq \frac{\tau_m^p}{1 + \gamma \tau_m^p} \quad (22)$$

With the above approximations, we can now define the following two scalar functions: $g_{\text{in}}(\cdot)$ and $g_{\text{out}}(\cdot)$ that are used in the GAMP algorithm. The input scalar function $g_{\text{in}}(\cdot)$ is simply defined as the posterior mean μ_n^x [22], i.e.

$$g_{\text{in}}(\hat{r}_n, \tau_n^r, \boldsymbol{\theta}) = \mu_n^x = \frac{\hat{r}_n}{1 + \eta_n \tau_n^r} \quad (23)$$

The scaled partial derivative of $\tau_n^r g_{\text{in}}(\hat{r}_n, \tau_n^r, \boldsymbol{\theta})$ with respect to \hat{r}_n is the posterior variance ϕ_n^x , i.e.

$$\tau_n^r \frac{\partial}{\partial \hat{r}_n} g_{\text{in}}(\hat{r}_n, \tau_n^r, \boldsymbol{\theta}) = \phi_n^x = \frac{\tau_n^r}{1 + \eta_n \tau_n^r} \quad (24)$$

The output scalar function $g_{\text{out}}(\cdot)$ is related to the posterior mean μ_m^z as follows

$$g_{\text{out}}(\hat{p}_m, \tau_m^p, \boldsymbol{\theta}) = \frac{1}{\tau_m^p} (\mu_m^z - \hat{p}_m) = \frac{1}{\tau_m^p} \left(\frac{\tau_m^p \gamma y_m + \hat{p}_m}{1 + \gamma \tau_m^p} - \hat{p}_m \right) \quad (25)$$

The partial derivative of $g_{\text{out}}(\hat{p}_m, \tau_m^p, \boldsymbol{\theta})$ is related to the posterior variance ϕ_m^z in the following way

$$\tau_m^p \frac{\partial}{\partial \hat{p}_m} g_{\text{out}}(\hat{p}_m, \tau_m^p, \boldsymbol{\theta}) = \frac{\phi_m^z - \tau_m^p}{\tau_m^p} = \frac{-\gamma \tau_m^p}{1 + \gamma \tau_m^p} \quad (26)$$

Given the above definitions of $g_{\text{in}}(\cdot)$ and $g_{\text{out}}(\cdot)$, the GAMP algorithm tailored to the considered sparse signal estimation problem can now be summarized as follows (details of the derivation of the GAMP algorithm can be found in [22]), in which a_{mn} denotes the (m, n) th entry of \mathbf{A} , $\mu_n^x(k)$ and $\phi_n^x(k)$ denote the posterior mean and variance of x_n at iteration k , respectively.

GAMP Algorithm

Initialization: given $\theta^{(t)}$; set $k = 0$, $\hat{s}_m^{(-1)} = 0, \forall m \in \{1, \dots, M\}$; $\{\mu_n^x(k)\}_{n=1}^N$ and $\{\phi_n^x(k)\}_{n=1}^N$ are initialized as the mean and variance of the prior distribution. Repeat the following steps until $\sum_n |\mu_n^x(k+1) - \mu_n^x(k)|^2 \leq \epsilon$, where ϵ is a pre-specified error tolerance.

Step 1. $\forall m \in \{1, \dots, M\}$:

$$\hat{z}_m(k) = \sum_n a_{mn} \mu_n^x(k)$$

$$\tau_m^p(k) = \sum_n a_{mn}^2 \phi_n^x(k)$$

$$\hat{p}_m(k) = \hat{z}_m(k) - \tau_m^p(k) \hat{s}_m(k-1)$$

Step 2. $\forall m \in \{1, \dots, M\}$:

$$\hat{s}_m(k) = g_{\text{out}}(\hat{p}_m(k), \tau_m^p(k), \theta^{(t)})$$

$$\tau_m^s(k) = -\frac{\partial}{\partial \hat{p}_m} g_{\text{out}}(\hat{p}_m(k), \tau_m^p(k), \theta^{(t)})$$

Step 3. $\forall n \in \{1, \dots, N\}$:

$$\tau_n^r(k) = \left(\sum_m a_{mn}^2 \tau_m^s(k) \right)^{-1}$$

$$\hat{r}_n(k) = \mu_n^x(k) + \tau_n^r(k) \sum_m a_{mn} \hat{s}_m(k)$$

Step 4. $\forall n \in \{1, \dots, N\}$:

$$\mu_n^x(k+1) = g_{\text{in}}(\hat{r}_n(k), \tau_n^r(k), \theta^{(t)})$$

$$\phi_n^x(k+1) = \tau_n^r(k) \frac{\partial}{\partial \hat{r}_n} g_{\text{in}}(\hat{r}_n(k), \tau_n^r(k), \theta^{(t)})$$

Output: $\{\hat{r}_n(k_0), \tau_n^r(k_0)\}$, $\{\hat{p}_m(k_0), \tau_m^p(k_0)\}$, and $\{\mu_n^x(k_0+1), \phi_n^x(k_0+1)\}$, where k_0 stands for the last iteration.

We have now derived an efficient algorithm to obtain approximate posterior distributions for the variables \mathbf{x} and $\mathbf{z} \triangleq \mathbf{Ax}$. We see that the GAMP algorithm no longer needs to compute the aforementioned matrix inverse. The dominating operations in each iteration is the simple matrix-vector multiplications, which scale as $\mathcal{O}(MN)$. Thus the computational complexity is significantly reduced. In the following, we discuss how to update the hyperparameters via the EM.

B. Hyperparameter Learning via EM

Given the current estimates of the hyperparameters $\theta^{(t)} \triangleq \{\alpha^{(t)}, \gamma^{(t)}\}$ and the observed data \mathbf{y} , the E-step computes the expected value (with respect to the hidden variables \mathbf{x}) of the complete log-posterior of $\{\alpha, \gamma\}$, i.e. $E_{\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)}}[\log p(\alpha, \gamma|\mathbf{x}, \mathbf{y})]$, where the operator $E_{\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)}}[\cdot]$ denotes the expectation with respect to the posterior distribution $p(\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)})$. This complete log-posterior is also referred to as the Q-function. The hyperparameters $\{\alpha, \gamma\}$ are then estimated by maximizing the Q-function, i.e.

$$\begin{aligned} \{\alpha^{(t+1)}, \gamma^{(t+1)}\} &= \arg \max_{\alpha, \gamma} Q(\alpha, \gamma|\alpha^{(t)}, \gamma^{(t)}) \\ &= \arg \max_{\alpha, \gamma} E_{\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)}}[\log p(\alpha, \gamma|\mathbf{x}, \mathbf{y})] \end{aligned} \quad (27)$$

Since

$$p(\alpha, \gamma|\mathbf{x}, \mathbf{y}) \propto p(\alpha)p(\mathbf{x}|\alpha)p(\gamma)p(\mathbf{y}|\mathbf{x}, \gamma) \quad (28)$$

the Q-function can be decomposed into a summation of two terms

$$\begin{aligned} Q(\alpha, \gamma|\alpha^{(t)}, \gamma^{(t)}) &= E_{\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)}}[\log p(\alpha)p(\mathbf{x}|\alpha)] \\ &\quad + E_{\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)}}[\log p(\gamma)p(\mathbf{y}|\mathbf{x}, \gamma)] \\ &\triangleq Q(\alpha|\alpha^{(t)}, \gamma^{(t)}) + Q(\gamma|\alpha^{(t)}, \gamma^{(t)}) \end{aligned} \quad (29)$$

We observe that in the above Q-function, the hyperparameters α and γ are separated from each other. This decouples the estimation of α and γ into two independent optimization problems. We first examine the update of α , i.e.

$$\alpha^{(t+1)} = \arg \max_{\alpha} Q(\alpha|\alpha^{(t)}, \gamma^{(t)}) \quad (30)$$

Recalling (11)–(13), $Q(\alpha|\alpha^{(t)}, \gamma^{(t)})$ can be expressed as

$$\begin{aligned} &Q(\alpha|\alpha^{(t)}, \gamma^{(t)}) \\ &\propto \sum_{n=1}^N \left((a-1) \log \alpha_n - b \alpha_n + \frac{1}{2} \log \eta_n - \frac{1}{2} \eta_n \langle x_n^2 \rangle \right) \end{aligned} \quad (31)$$

where $\langle x_n^2 \rangle$ denotes the expectation with respect to the posterior distribution $p(\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)})$. Here we use $\hat{p}(x_n|\mathbf{y}, \hat{r}_n(k_0), \tau_n^r(k_0), \theta^{(t)})$, i.e. the approximate posterior distribution of x_n obtained from the GAMP algorithm, to replace the true posterior distribution $p(\mathbf{x}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)})$ in computing the expectation. Since $\hat{p}(x_n|\mathbf{y}, \hat{r}_n(k_0), \tau_n^r(k_0), \theta^{(t)})$ follows a Gaussian distribution with its mean and variance given by (18)–(19), we have

$$\langle x_n^2 \rangle = \frac{(\hat{r}_n(k_0))^2}{(1 + \eta_n^{(t)} \tau_n^r(k_0))^2} + \frac{\tau_n^r(k_0)}{1 + \eta_n^{(t)} \tau_n^r(k_0)} \quad (32)$$

We see that in the Q-function (31), hyperparameters are entangled with each other due to the logarithm term $\log \eta_n$ (note that η_n , defined in (12), is a function of α). In this case, an analytical solution to the optimization (30) is difficult to obtain. Gradient descent methods can certainly be used to search for the optimal solution. Nevertheless, for gradient descent methods, there is no explicit formula for the hyperparameter update. Also, gradient-based methods involve higher computational complexity as compared with an analytical update rule. Here we consider an alternative strategy which aims at finding a simple, analytical sub-optimal solution of (30). Such an analytical sub-optimal solution can be obtained by examining the optimality condition of (30). Suppose $\alpha^* = \{\alpha_n^*\}$ is the optimal solution of (30). Then the first derivative of the Q-function with respect to α equals to zero at the optimal point, i.e.

$$\frac{\partial Q(\alpha|\theta^{(t)})}{\partial \alpha_n} \Big|_{\alpha=\alpha^*} = \frac{a-1}{\alpha_n^*} - b + \frac{1}{2} \nu_n^* - \frac{1}{2} \omega_n = 0 \quad (33)$$

where

$$\nu_n^* \triangleq \frac{1}{\eta_n^*} + \beta \sum_{i \in N_{(n)}} \frac{1}{\eta_i^*} \quad (34)$$

$$\omega_n \triangleq \langle x_n^2 \rangle + \beta \sum_{i \in N_{(n)}} \langle x_i^2 \rangle \quad (35)$$

Since all hyperparameters $\{\alpha_n\}$ and β are non-negative, it can be easily verified $(1/\alpha_n^*) > (1/\eta_n^*) > 0$, and $(1/\beta\alpha_n^*) > (1/\eta_i^*) > 0$ for $i \in N_{(n)}$. Therefore ν_n^* is bounded by

$$\frac{5}{\alpha_n^*} > \nu_n^* > 0 \quad (36)$$

Consequently we have

$$\frac{a+1.5}{\alpha_n^*} > \frac{a-1}{\alpha_n^*} + \frac{1}{2}\nu_n^* > \frac{a-1}{\alpha_n^*} \quad (37)$$

Combining (33) and (37), we reach that α_n^* is within the range

$$\alpha_n^* \in \left[\frac{a-1}{0.5\omega_n+b}, \frac{a+1.5}{0.5\omega_n+b} \right] \quad \forall n \quad (38)$$

A sub-optimal solution to (30) can therefore simply be chosen as

$$\alpha_n^{(t+1)} = \frac{a-1}{0.5\omega_n+b} \quad \forall n \quad (39)$$

We see that the solution (39) provides a simple rule for the update of α . Also, notice that the update rule (39) resembles that of the conventional sparse Bayesian learning work [4] except that ω_n is equal to $\langle x_n^2 \rangle$ for the conventional sparse Bayesian learning method, while for our case, ω_n is a weighted summation of $\langle x_n^2 \rangle$ and $\langle x_i^2 \rangle$ for $i \in N_{(n)}$. Numerical results show that for a properly chosen value of a , this update rule, although sub-optimal, guarantees an exact recovery and provides superior recovery performance. Specifically, $a = 1.5$ is empirically proven to be a robust choice which ensures that the proposed algorithm delivers stable recovery performance.

We now discuss the update of the hyperparameter γ , the inverse of the noise variance. Since the GAMP algorithm also provides an approximate posterior distribution for the noiseless output $z = \mathbf{A}\mathbf{x}$, we can simply treat z as hidden variables when learning the noise variance, i.e.

$$\begin{aligned} \gamma^{(t+1)} &= \arg \max_{\gamma} E_{\mathbf{z}|\mathbf{y}, \alpha^{(t)}, \gamma^{(t)}} [\log p(\gamma)p(\mathbf{y}|\mathbf{z}, \gamma)] \\ &\triangleq \arg \max_{\gamma} Q(\gamma|\alpha^{(t)}, \gamma^{(t)}) \end{aligned} \quad (40)$$

Taking the partial derivative of the Q-function with respect to γ gives

$$\frac{\partial Q(\gamma|\alpha^{(t)}, \gamma^{(t)})}{\partial \gamma} = \frac{c-1}{\gamma} - d + \frac{M}{2\gamma} - \frac{1}{2} \sum_{m=1}^M \langle (y_m - z_m)^2 \rangle \quad (41)$$

where $\langle \cdot \rangle$ denotes the expectation with respect to $p(z_m|\mathbf{y}, \hat{p}_m(k_0), \tau_m^p(k_0), \boldsymbol{\theta}^{(t)})$, i.e. the approximate posterior distribution of z_m . Recalling that the approximate posterior of z_m follows a Gaussian distribution with its mean and variance given by (21)–(22), we have

$$\langle (y_m - z_m)^2 \rangle = (y_m - \mu_m^z)^2 + \phi_m^z \quad (42)$$

where μ_m^z and ϕ_m^z are given by (21)–(22), with $\{\hat{p}_m, \tau_m^p\}$ replaced by $\{\hat{p}_m(k_0), \tau_m^p(k_0)\}$, and γ replaced by $\gamma^{(t)}$. Setting the derivative equal to zero, we obtain the update rule for γ as

$$\gamma^{(t+1)} = \frac{M+2c-2}{2d + \sum_m \langle (y_m - z_m)^2 \rangle} \quad (43)$$

So far we have completed the development of our GAMP-based pattern-coupled sparse Bayesian learning algorithm. For clarify, we now summarize our proposed PCSBL-GAMP algorithm as follows.

PCSBL-GAMP Algorithm

1. Initialization: given $\alpha^{(0)}$ and $\gamma^{(0)}$.
2. For $t \geq 0$: given $\alpha^{(t)}$ and $\gamma^{(t)}$, call the GAMP algorithm. Based on the outputs of the GAMP algorithm, update the hyperparameters $\alpha^{(t+1)}$ and $\gamma^{(t+1)}$ according to (39) and (43).
3. Continue the above iteration until the difference between two consecutive estimates of \mathbf{x} is negligible, or a maximum number of iterations is reached.

IV. SIMULATION RESULTS

We now carry out experiments to illustrate the performance of our proposed algorithm, also referred to as PCSBL-GAMP algorithm, and its comparison with other existing methods. The performance of the proposed algorithm³ will be examined using both synthetic and real data. The parameters a, b, c, d for our proposed algorithm are set equal to $a = 1.5$, $b = 10^{-6}$, $c = 1$, and $d = 10^{-6}$ throughout our experiments. The relevance parameter β is set equal to $\beta = 1$ in our experiments. In fact, our empirical results suggest that its choice is not very critical to the recovery performance as long as β is set into the region $\beta \in [0.1, 1]$.

A. Synthetic Data

We first evaluate the recovery performance of the PCSBL-GAMP method using the synthetic data. In our simulations, we generate a one-dimensional block-sparse signal in a way similar to [21]. Suppose the N -dimensional sparse signal contains K nonzero coefficients (K is also denoted as the sparsity level) which are partitioned into T blocks with random sizes and random locations. The nonzero coefficients of the sparse signal \mathbf{x} and the measurement matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ are randomly generated with each entry independently drawn from a normal distribution, and then the sparse signal \mathbf{x} and columns of \mathbf{A} are normalized to unit norm. Fig. 1 depicts the success rate of the proposed PCSBL-GAMP method vs. the ratio M/N , where we set $N = 200$, $K = 40$ and $T = 6$. The success rate is computed as the ratio of the number of successful trials to the total number of independent runs. A trial is considered successful if the normalized squared recovery error $\|\mathbf{x} - \hat{\mathbf{x}}\|^2 / \|\mathbf{x}\|^2$ is no greater than 10^{-6} , where $\hat{\mathbf{x}}$ denotes the estimate of the sparse signal. The success rates of the EM-based PC-SBL method [21] (referred to as the PCSBL-EM), the conventional SBL [4], and the basis pursuit (BP) method [2], [24] are also included for comparison. From Fig. 1, we see that the PCSBL-GAMP method achieves almost the same performance as that of the PCSBL-EM method, and presents a significant performance advantage over the SBL and BP methods due to exploiting the underlying block-sparse structures. The average run times of respective algorithms are

³Matlab codes for our algorithm are available at <http://www.junfang-uestc.net/codes/PCSBL-GAMP.rar>

a function of the signal dimension N is plotted in Fig. 2, where we set $M = 0.4N$. Results are averaged over 100 independent runs. We see that the PCSBL-GAMP requires much less run time than the PCSBL-EM, particularly when the signal dimension N is large. Also, it can be observed that the average run time of the PCSBL-EM grows rapidly with an increasing N , whereas the average run time of the PCSBL-GAMP increases very slowly. This observation coincides with our computational complexity analysis.




Fig. 1. Success rates vs. the ratio M/N .




Fig. 2. Average run times vs. N .

We also carry out experiments using patches of letters “C” and “S” (16×16 pixels) with black background, where most of the pixels on the patches are zeros and the nonzero coefficients exhibit irregular block patterns. White Gaussian noise is added to the patches. We compare our method with other block-sparse signal recovery algorithms, namely, the cluster-structured MCMC algorithm (Cluss-MCMC) [17], the Boltzman machine-based greedy pursuit algorithm (BM-MAP-OMP) [18], and the block sparse Bayesian learning method (BSBL) method [15], [20]. Note that although the BM-MAP-OMP and the BSBL algorithms are developed for one-dimensional sparse signal recovery, we extend their methods to the two-dimensional scenario. In our simulations, model

parameters used by the competing algorithms are adjusted to achieve the best performance. For the BSBL algorithms, the block size parameter h is set equal to 2. Fig. 3 depicts the original patches and the patches reconstructed by respective algorithms, where we set $M = 80$, and the signal-to-noise ratio (SNR) is set to 20dB with $\text{SNR} \triangleq 10 \log(\|\mathbf{Ax}\|_2^2 / M\sigma^2)$. It can be seen that our proposed PCSBL-GAMP method provides the best visual quality with recognizable letters, whereas the letters reconstructed by other algorithms have considerably lower quality, particularly for the Cluss-MCMC and the BM-MAP-OMP methods. This result also implies that our proposed method is flexible to accommodate any irregular cluster patterns.

B. Satellite Image Recovery

In this subsection, we carry out experiments on a non-negative 256×256 satellite image⁴. The image is sparse in the spatial domain, with only 6678 (approximately 10.2% of total pixels) nonzero pixels. In our experiments, compressive measurements are corrupted by additive i.i.d. Gaussian noise, i.e. $\mathbf{y} = \mathbf{Ax} + \mathbf{w}$, where the image is represented as a one-dimensional vector \mathbf{x} . The sensing matrix \mathbf{A} is chosen to be the same as that used in [25], i.e. $\mathbf{A} = \mathbf{\Phi}\mathbf{\Psi}\mathbf{S}$, where $\mathbf{\Phi} \in \{0, 1\}^{M \times N}$ and its rows are randomly selected from the $N \times N$ identity matrix, $\mathbf{\Psi} \in \{-1, 1\}^{N \times N}$ is a Hadamard transform matrix, $\mathbf{S} \in \mathbb{R}^{N \times N}$ is a diagonal matrix with its entries randomly chosen from $\{-1, 1\}$. Sensing using such a measurement matrix can be executed using a fast binary algorithm, which makes the hardware implementation simple. Note that the BM-MAP-OMP, the BSBL and the Cluss-MCMC methods were not included in this experiment due to their prohibitive computational complexity when the signal dimension is large. Instead, we compare our method with some other computationally efficient GAMP-based methods, namely, the EM-NNGM-GAMP method [25] and the EM-GM-AMP method [26]. These two methods have demonstrated state-of-the-art recovery performance in a series of experiments. The spectral projected gradient (SPG) method (referred to as SPGL1) which was developed in [27] to efficiently solve the basis pursuit or basis pursuit denoising optimizations is also included for comparison. Among these methods, only the EM-NNGM-GAMP algorithm exploits the non-negativity of the satellite image. Therefore the signals recovered by these algorithms, except the EM-NNGM-GAMP, may contain negative coefficients. These negative coefficients are manually set to zero in our simulations. Fig. 4 shows the original satellite image and the images reconstructed by respective algorithms, where we set $M = 0.15N$ and $\text{SNR} = 60\text{dB}$. It can be seen that the PCSBL-GAMP offers a significantly better image quality as compared with other methods. Fig. 5 plots the normalized mean square errors (NMSEs) of respective algorithms vs. the ratio M/N . Results are averaged over 100 independent trials, with the sensing matrix \mathbf{A} randomly generated for each trial. The recovered negative coefficients are kept unaltered in calculating the NMSEs. We see that our proposed PCSBL-GAMP method outperforms other algorithms by a big margin

⁴Image data is available at <http://sourceforge.net/projects/gampmatlab>.




Fig. 3. Original patches of letters and patches reconstructed by respective algorithms.




Fig. 4. Original satellite image and images reconstructed by respective algorithms.

for a small ratio M/N (e.g. $M/N \leq 0.25$), where data acquisition is practically appealing due to high compression rates.




Fig. 5. Satellite image: NMSEs vs. the ratio M/N .

C. Background Subtraction

Background subtraction, also known as foreground detection, is a technique used to automatically detect and track moving objects in videos from static cameras. Usually, the

foreground innovations are sparse in the spatial image domain. By exploiting this sparsity, the sparse foreground innovations within a scene can be reconstructed using compressed measurements, which relieves the communications burden placed on data transmission [8]. Specifically, the idea is to reconstruct the foreground image from the difference between the compressed measurements of the background image and the compressed measurements of the test image [8]

$$\begin{aligned} \mathbf{y}_f &= \mathbf{y}_t - \mathbf{y}_b = \mathbf{A}(\mathbf{x}_t - \mathbf{x}_b) \\ &= \mathbf{A}\mathbf{x}_f \end{aligned}$$

where \mathbf{x}_t and \mathbf{x}_b represent the test and the background images, respectively; \mathbf{y}_t and \mathbf{y}_b denote the compressed measurements of the test and background images, respectively; and \mathbf{x}_f is the foreground image to be recovered. In our experiments, we use the Convoy2 data set that was used in [28]. The Convoy2 data set was collected on the Spesutie island, consisting of a video sequence with 260 frames and one background frame recorded by a single static camera. The video sequence has a dynamic sparse foreground as vehicles enter and exit the filed of view over time. We first choose the 40th frame of the Convoy2 data set as a test image, which is shown in Fig. 6. The background image and the foreground image are also included in Fig. 6. The foreground image is regarded as the groundtruth image. This foreground image, however, does not have a pure background since $\mathbf{x}_f = \mathbf{x}_t - \mathbf{x}_b$ is not an exactly sparse

signal and contains many small nonzero components. In our experiments, the original images of 480×381 pixels are resized to 120×96 pixels. For the resized foreground image, we have a total number of 923 coefficients whose magnitudes are greater than 10^{-2} , thus the percentage of nonzero coefficients is $923/(120 \times 96) = 8.01\%$. Again, the BM-MAP-OMP, BSBL, and the Cluss-MCMC methods are not included due to their prohibitive computational complexity. Here we compare our method with the EM-GM-AMP method [26] and the EM-BG-AMP method [29]. Fig. 7 depicts images reconstructed by respective algorithms, where we use only $M = 0.1N$ measurements. The measurement matrix \mathbf{A} is randomly generated with each entry independently drawn from a normal distribution. We see that our proposed PCSBL-GAMP method provides the finest image quality with a clear appearance of the vehicle, whereas the object silhouettes recovered by other methods are hardly recognizable. In our next experiments, frames from the 10th to 60th are used as test images. For each test image, we use $M = 0.1N$ measurements to recover the difference (foreground) image. Fig. 8 shows the NMSEs of respective algorithms vs. the frame number, where the NMSEs are obtained by averaging over 100 independent runs, with the measurement matrix \mathbf{A} randomly generated for each run. From Fig. 8, we observe that our proposed method presents a significant performance advantage over other methods.

V. CONCLUSIONS

We developed a pattern-coupled sparse Bayesian learning method for recovery of two-dimensional block-sparse signals whose cluster patterns are unknown *a priori*. A two-dimensional pattern-coupled hierarchical Gaussian prior model is introduced to characterize and exploit the pattern dependencies among neighboring coefficients. The proposed pattern-coupled hierarchical model is effective and flexible to capture any underlying block-sparse structures, without requiring the prior knowledge of the block partition. An expectation-maximization (EM) strategy is employed to infer the maximum a posterior (MAP) estimate of the hyperparameters, along with the posterior distribution of the sparse signal. Additionally, the generalized approximate message passing (GAMP) algorithm is embedded in the EM framework to efficiently compute an approximation of the posterior distribution of hidden variables, which results in a significant reduction in computational complexity. Numerical results show that our proposed algorithm presents a substantial performance advantage over other existing state-of-the-art methods in image recovery.

REFERENCES

- [1] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Information Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.
- [2] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [3] D. Wipf and S. Nagarajan, "Iterative reweighted ℓ_1 and ℓ_2 methods for finding sparse solutions," *IEEE Journals of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 317–329, Apr. 2010.
- [4] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.




Fig. 8. NMSEs vs. the frame number.

- [5] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Signal Processing*, vol. 56, no. 6, pp. 2346–2356, June 2008.
- [6] Z. Yang, L. Xie, and C. Zhang, "Off-grid direction of arrival estimation using sparse Bayesian inference," *IEEE Trans. Signal Processing*, vol. 61, no. 1, pp. 38–42, Jan. 2013.
- [7] L. Wang, L. Zhao, G. Bi, C. Wan, and L. Yang, "Enhanced ISAR imaging by exploiting the continuity of the target scene," *IEEE Trans. Geoscience and Remote Sensing*, no. 9, pp. 5736–5750, Sept. 2014.
- [8] V. Cevher, A. Sankaranarayanan, M. F. Duarte, D. Reddy, R. G. Baraniuk, and R. Chellappa, "Compressive sensing for background subtraction," in *European Conf. Comp. Vision (ECCV)*, Marseille, France, October 12–18 2008.
- [9] N. Vaswani and W. Lu, "Modified-CS: modifying compressive sensing for problems with partially known support," *IEEE Trans. Signal Processing*, no. 9, pp. 4595–4607, September 2010.
- [10] Y. C. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Information Theory*, vol. 55, no. 11, pp. 5302–5316, Nov. 2009.
- [11] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Information Theory*, vol. 56, no. 4, pp. 1982–2001, Apr. 2010.
- [12] Y. C. Eldar, P. Kuppinger, and H. Bölcskei, "Block-sparse signals: uncertainty relations and efficient recovery," *IEEE Trans. Information Theory*, vol. 58, no. 6, pp. 3042–3054, June 2010.
- [13] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. R. Statist. Soc. B*, vol. 68, pp. 49–67, 2006.
- [14] D. P. Wipf and B. D. Rao, "An empirical Bayesian strategy for solving the simultaneous sparse approximation problem," *IEEE Trans. Signal Processing*, vol. 55, no. 7, pp. 3704–3716, July 2007.
- [15] Z. Zhang and B. D. Rao, "Sparse signal recovery with temporally correlated source vectors using sparse Bayesian learning," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 5, pp. 912–926, Sept. 2011.
- [16] L. He and L. Carin, "Exploiting structure in wavelet-based Bayesian compressive sensing," *IEEE Trans. Signal Processing*, vol. 57, no. 9, pp. 3488–3497, Sept. 2009.
- [17] L. Yu, H. Sun, J. P. Barbot, and G. Zheng, "Bayesian compressive sensing for cluster structured sparse signals," *Signal Processing*, vol. 92, pp. 259–269, 2012.
- [18] T. Peleg, Y. C. Eldar, and M. Elad, "Exploiting statistical dependencies in sparse representations for signal recovery," *IEEE Trans. Signal Processing*, vol. 60, no. 5, pp. 2286–2303, May 2012.
- [19] A. Drămeau, C. Herzet, and L. Daudet, "Boltzmann machine and mean-field approximation for structured sparse decompositions," *IEEE Trans. Signal Processing*, vol. 60, no. 7, pp. 3425–3438, July 2012.
- [20] Z. Zhang and B. D. Rao, "Extension of SBL algorithms for the recovery of block sparse signals with intra-block correlation," *IEEE Trans. Signal Processing*, vol. 61, no. 8, pp. 2009–2015, Apr. 2013.
- [21] J. Fang, Y. Shen, H. Li, and P. Wang, "Pattern-coupled sparse Bayesian learning for recovery of block-sparse signals," *IEEE Trans. Signal Processing*, no. 2, pp. 360–372, Jan. 2015.
- [22] S. Rangan, "Generalized approximate message passing for estimation



Fig. 6. From left to right: the test image \mathbf{x}_t (40th frame of the Convoy2 data set), the background image \mathbf{x}_b , the foreground image \mathbf{x}_f .




Fig. 7. Foreground images reconstructed by respective algorithms.

- with random linear mixing,” in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Saint Petersburg, Russia, Aug. 2011.
- [23] D. L. Donoho, A. Maleki, and A. Montanari, “Message passing algorithms for compressed sensing: I. motivation and construction,” in *Proc. Inf. Theory Workshop*, Cairo, Egypt, Jan. 2010.
 - [24] E. Candés and T. Tao, “Decoding by linear programming,” *IEEE Trans. Information Theory*, no. 12, pp. 4203–4215, Dec. 2005.
 - [25] J. P. Vila and P. Schniter, “An empiricla-Bayes approach to recovering linearly constrained non-negative sparse signals,” *IEEE Trans. Signal Processing*, vol. 62, no. 18, pp. 4689–4703, Sep 2014.
 - [26] —, “Expectation-maximization Gaussian-mixture approximate message passing,” *IEEE Trans. Signal Processing*, no. 19, pp. 4658–4672, Oct. 2013.
 - [27] E. van den Berg and M. P. Friedlander, “Probing the Pareto frontier for basis pursuit solutions,” *SIAM J. Sci. Comput.*, vol. 31, no. 2, pp. 890–912, 2008.
 - [28] G. Warnell, D. Reddy, and R. Chellappa, “Adaptive rate compressive sensing for background subtraction,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Kyoto, Japan, March 25-30 2012.
 - [29] J. P. Vila and P. Schniter, “Expectation-maximization Bernoulli-Gaussian approximate message passing,” in *45th Asilomar Conference on Signals, Syntems and computers*, Pacific Grove, CA, USA, November 6-9 2011.